



US 20140157274A1

(19) **United States**(12) **Patent Application Publication**
Ballani et al.(10) **Pub. No.: US 2014/0157274 A1**(43) **Pub. Date: Jun. 5, 2014**(54) **OFFERING NETWORK PERFORMANCE
GUARANTEES IN MULTI-TENANT
DATACENTERS**(52) **U.S. Cl.**
CPC **G06F 9/45533** (2013.01)
USPC **718/1**(71) Applicant: **Microsoft Corporation**, Redmond, WA
(US)(72) Inventors: **Hitesh Ballani**, Cambridge (GB); **Paolo
Costa**, Cambridge (GB); **Thomas
Karagiannis**, Cambridge (GB); **Antony
Rowstron**, Cambridge (GB)(73) Assignee: **Microsoft Corporation**, Redmond, WA
(US)(21) Appl. No.: **14/177,202**(22) Filed: **Feb. 10, 2014****Related U.S. Application Data**(63) Continuation of application No. 13/176,901, filed on
Jul. 6, 2011, now Pat. No. 8,671,407.**Publication Classification**(51) **Int. Cl.**
G06F 9/455 (2006.01)(57) **ABSTRACT**

Methods of offering network performance guarantees in multi-tenant datacenters are described. In an embodiment, a request for resources received at a datacenter from a tenant comprises a number of virtual machines and a performance requirement, such as a bandwidth requirement, specified by the tenant. A network manager within the datacenter maps the request onto the datacenter topology and allocates virtual machines within the datacenter based on the available slots for virtual machines within the topology and such that the performance requirement is satisfied. Following allocation, stored residual capacity values for elements within the topology are updated according to the new allocation and this updated stored data is used in mapping subsequent requests onto the datacenter. The allocated virtual machines form part of a virtual network within the datacenter which is allocated in response to the request and two virtual network abstractions are described: virtual clusters and virtual oversubscribed clusters.

